Server side processing of petascale data sets from the Global Cloud Resolving Model

K. Schuchardt[1], B. Palmer[1], T. Elsethagen[1], J Daily[1], and A Koontz[1]
[1]Pacific Northwest National Laboratory

Summary

The Global Cloud Resolving Model being developed at Colorado State University under Professor David Randall, will simulate the entire globe at a resolution of 2-4 km. At such scales, a single snapshot of data will result in a terabyte or more of data. Even early simulations of modest time scales will result in petabytes of data. Instead of moving data to local workstations to perform analysis as is commonly done today, it is critical to move services to the data and take advantage of parallel computers and file systems. We are developing such services both for data reduction and visualization of large scale data sets based on the geodesic grid. These services will be integrated with the Earth Systems Grid (ESG), taking advantage of the infrastructure already provided by ESG while driving requirements for supporting other high-resolution model outputs.

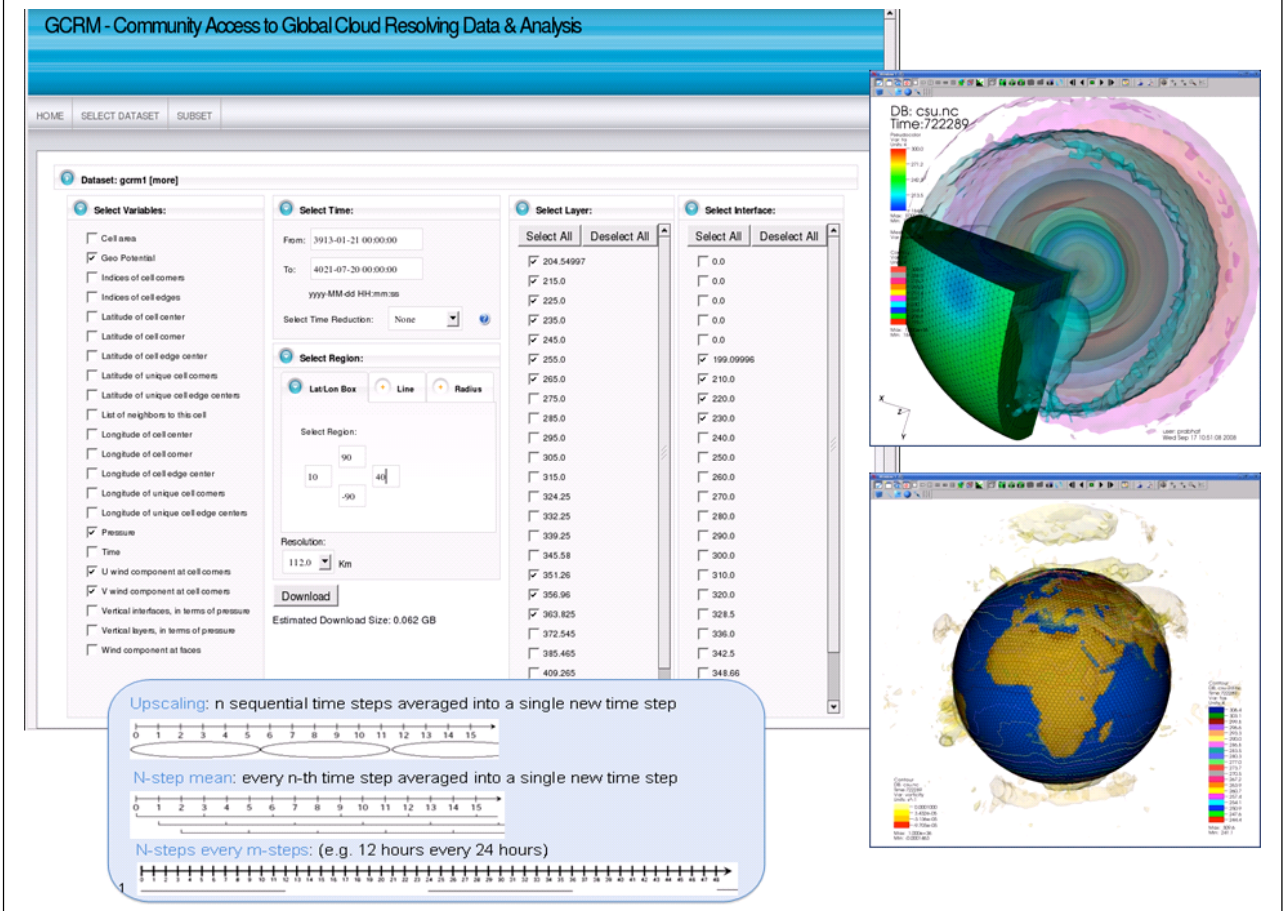

Figure 1. Screenshot of prototype portal that shows data subsetting aand data reduction capabilities for GCRM data. This example shows relatively low resolution data, but conveys the key requirement of greatly reducing data set sizes through sampling, coarsening, and averaging techniques, such that they may more readily be transferred for analysis. Also shown, is the concept of generating visualizations and plots (not shown) as a server side function. While ESG currently supports this type of analysis, new high performance tools will be required to support GCRM data. The visualizations where provided by Prabhat at Lawrence Berkeley National Laboratory.

The Global Cloud Resolving Model (GCRM) being developed at Colorado State University under Professor David Randall, will simulate the entire globe at a resolution of 2-4 km. At such scales, the model can crudely represent individual large clouds, an important step forward as the representation of clouds in weather prediction and general circulation models has been identified as one of the greatest uncertainties in predictive climate modeling. The GCRM will also simulate dynamic interactions across a wide range of spatial scales including squall lines, tropical cyclones, midlatitude baroclinic waves, large planetary waves, and monsoons.

At 2-4 km resolution, approximately 100 vertical levels will be required. At 4km, the horizontal grid will contain 41 million cells, 82 million corners, and 120 million edges. Data is modeled on each of these grid locations. Assuming a mix of 30 3D variables, a single data snapshot will result in a terabyte or more of output. A one month simulation, saving hourly snapshots, will result in hundreds of terabytes on up to a petabyte. A full years' simulation under these parameters will result in many petabytes of data. In the near term, the model will be run on these time scales rather than the tens to hundreds of years typical of climate simulations due to the intense computational requirements.

The GCRM represents an extreme output case in the development of high resolution models. However, many other climate simulators are moving toward higher resolutions and very large data sets. This shift necessitates a move to parallel programs for data manipulation and analysis as well as a conceptual shift of deploying services with the data rather than moving data to services. These services can take the form of data reduction techniques that reduce data sets to sizes suitable for download as well as visualizations and analyses generated on the server side. We are currently developing these services specifically for the GCRM, in the form of custom parallel data manipulation tools and extensions to the VisIt visualization system, and will integrate them with the Earth Systems Grid (ESG) to take advantage of important ESG capabilities such as flexible searching and browsing, authentication and fine grained access controls, seamless access to off-line data, and alternatives for high throughput data downloads.

Integration of GCRM data into the ESG framework will certainly reveal new challenges and requirements. Through ongoing discussions, we have identified several key areas where we expect to focus our efforts. A crucial first step is to ensure that the metadata model is sufficiently rich to support the geodesic grid and possibly semi-structured or unstructured grids in general. The grid also impacts the features that may be required for data reduction, such as the specification of averaging or coarsening algorithms, which in turn, is likely to impact both the user interface design and the design of the protocol between the gateway and data nodes to allow alternative data node service providers and operations. Because the data sets will be large, most operations will require asynchronous handling of user requests and integration with a queue system to manage the parallel analysis requests. This requires working closely with computing centers to address policy issues related to running such services as well as managing disk quotas. Ultimately, we expect to take advantage of ESG to not only provide basic GCRM data services to a broader community, but also to support GCRM model comparison.

For further information on this subject contact:

- Name: Karen Schuchardt
  Organization: Pacific Northwest National Laboratory
  Email: karen.schuchardt@pnl.gov

- Name: Dean N. Williams.
  Organization: Lawrence Livermore National Laboratory
  Email: williams13@llnl.gov

This research is supported by the U. S. Department of Energy's Office of Science under the Scientific Discovery through Advanced Computing (SciDAC) program under contract number DE-AC05-76RLO1830.